REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
22-08-2014	Conference Proceeding	-

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
An Improved Evolution-COnstructed (iECO) Features Framework	W911NF-14-1-0114
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHORS	5d. PROJECT NUMBER
Stanton Price, Derek Anderson, Robert Luke	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
Mississippi State University PO Box 6156 Mississippi State, MS 39762 -6156	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 65266-CS.2

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

In image processing and computer vision, significant progress has been made in feature learning for exploiting important cues in data that elude non-learned features. While the field of deep learning has demonstrated state-of-the-art performance, the Evolution-COnstructed (ECO) work of Lillywhite et. al has the advantage of interpretability, and it does not pre-dispose the solution to one of convolution. This paper presents a novel approach for extending the ECO framework. We achieve this through two overarching ideas. First, we address a potential major shortcoming of ECO features — the 'features' themselves. The so-called ECO features are simply a

15. SUBJECT TERMS
evolution-constructed features, ECO, feature learning, infrared, explosive hazard detection

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Derek Anderson
UU	UU	UU	UU		19b. TELEPHONE NUMBER 662-325-3530

Report Title

An Improved Evolution-COnstructed (iECO) Features Framework

ABSTRACT

In image processing and computer vision, significant progress has been made in feature learning for exploiting important cues in data that elude non-learned features. While the field of deep learning has demonstrated state-of-the-art performance, the Evolution-COnstructed (ECO) work of Lillywhite et. al has the advantage of interpretability, and it does not pre-dispose the solution to one of convolution. This paper presents a novel approach for extending the ECO framework. We achieve this through two overarching ideas. First, we address a potential major shortcoming of ECO features-- the 'features' themselves. The so-called ECO features are simply a transformed image that has been unrolled into a large one dimensional vector. We propose employing feature descriptors to extract pertinent information from the ECO imagery. Furthermore, it is our hypothesis that there exists a unique set of transforms for each feature descriptor used on a given problem domain that leads to the descriptors extracting maximal discriminative information. Second, we introduce constraints on each individual's chromosome to promote population diversity and prevent infeasible solutions. We show through experiments that our proposed iECO framework results in, and benefits from, a unique series of transforms for each descriptor being learned and maintaining population diversity.

Conference Name: SSCI

Conference Date: December 20, 2014

An Improved Evolution-COnstructed (iECO) Features Framework

Distribution Statement A: Approved for public release; distribution is unlimited.

Stanton R. Price and Derek T. Anderson
Electrical and Computer Engineering Department
Mississippi State University
Mississippi State, MS
Email: sp438@msstate.edu, anderson@ece.msstate.edu

Robert H. Luke
U.S. Army RDECOM CERDEC
Night Vision and Electronic Sensors Directorate
Fort Belvoir, VA
Email: info@nvl.army.mil

Abstract—In image processing and computer vision, significant progress has been made in feature learning for exploiting important cues in data that elude non-learned features. While the field of deep learning has demonstrated state-of-the-art performance, the Evolution-COnstructed (ECO) work of Lillywhite et. al has the advantage of interpretability, and it does not predispose the solution to one of convolution. This paper presents a novel approach for extending the ECO framework. We achieve this through two overarching ideas. First, we address a potential major shortcoming of ECO features– the "features" themselves. The so-called ECO features are simply a transformed image that has been unrolled into a large one dimensional vector. We propose employing feature descriptors to extract pertinent information from the ECO imagery. Furthermore, it is our hypothesis that there exists a unique set of transforms for each feature descriptor used on a given problem domain that leads to the descriptors extracting maximal discriminative information. Second, we introduce constraints on each individual's chromosome to promote population diversity and prevent infeasible solutions. We show through experiments that our proposed iECO framework results in, and benefits from, a unique series of transforms for each descriptor being learned and maintaining population diversity.

I. INTRODUCTION

Object recognition is a very challenging task, but one that has garnered increased interest in recent years. In 2013, Lillywhite et. al [1] put forth the Evolution-COnstructed (ECO) features framework for general object recognition. Their approach, while novel and shown to produce state-of-the-art results, left a number of areas asking for further investigation and improvement. Herein, we present a novel approach to address two major shortcomings of the ECO features framework. Specifically, speaking in high-level terms, we employ feature descriptors and introduce diversity promoting constraints to the genetic algorithm (GA).

Feature descriptors are commonly used in the computer vision and image processing fields to robustly capture pertinent information for object recognition. However, they are typically taken at face-value. That is, for a given recognition task, a feature descriptor, or more commonly, a set of feature descriptors are extracted from the original imagery. In a sense, it is assumed that, if using the proper feature descriptor(s) for a given task and there is some amount of class separability, the feature descriptor(s) will adequately capture this information.

It is then the job of the classifier, e.g., Minimum Risk Bayes Decision Theoretic Classifier, k-nearest neighbor (kNN), support vector machine (SVM), etc., to find the discriminatory information for proper classification. Herein, we propose a novel ideology for approaching feature descriptors in a non-traditional way.

The ECO features framework uses a standard GA to learn a series of transforms that leads to the discovery of their so-called ECO features (a more thorough description is given in Section II). There are no direct mechanisms incorporated into the GA, outside of mutation, that have been put in place to promote diversity within the population. This is a major potential shortcoming. For one, this framework has a incredibly massive search space. Referring to [1], a total of 27 image transforms where available, with all but 6 having at least one parameter associated with it (one had as many as 6!). Additionally, there are four parameters used to select sub-regions in which the transforms were to be performed on. It is rather easy to recognize that this search space is incredibly large and very likely to get stuck searching only a relatively small portion of the space if no specialized diversity promoting method is implemented. Thus, we introduce constraints on each individual's chromosome to ensure population diversity.

We do acknowledge that in [1], speciation was implemented with the intention of obtaining a diverse solution. Therein, speciation was incorporated by learning ECO features in multiple small population sizes rather than using one large population. For further details on their implementation, we refer the reader to [1].

The remainder of this paper is organized as follows. In Section II, we put forth our improved ECO (iECO) framework. Specifically, we present a detailed discussion of our approach to using feature descriptors in Section II-A, and we formally introduce the diversity promoting constraints for the GA in Section II-B. An analysis of our proposed iECO framework is presented in Section III. Finally, we conclude this paper and provide insight to our future work in Section IV.

II. iECO FRAMEWORK

We begin this section with a brief overview of the ECO features framework (for a detailed discussion, we refer the

reader to [1]). The ECO features framework is an interesting method for approaching object recognition. It attempts to fully automate the process of feature construction (to an extent– if we are being sensitive, we could argue that potential solutions were imposed through the selection of transforms made available to the learner in addition to the assumption that this is the best way to recognize discriminative information from imagery). This is achieved using a GA to learn a series of image transforms and the region in the imagery that is best suited for such an ordering of transforms. In regards to our work, the image transforms that are available to the GA, their number of free parameters, and their gene identifier is shown in Table I.

TABLE I.   LIST OF IMAGE TRANSFORMS AVAILABLE TO THE GA, THEIR NUMBER OF FREE PARAMETERS, AND A GENE IDENTIFIER FOR EACH TRANSFORM.

Gene ID	Image Transform	$|\phi|$
0	Harris Corner Detector	1
1	Gradient	1
2	Square Root	0
3	Gaussian Blur	1
4	Hough Circle	1
5	Median Blur	1
6	Canny	0
7	Rank Transform	0
8	Log	0
9	Sobel	0
10	Difference of Gaussian	2
11	Erode	1
12	Dilate	1
13	Contrast Limited Adaptive Histogram Equalization	4
14	Distance Transform	0
15	Histogram Equalization	0
16	Laplacian Edge	1
17	Maximally Stable Extremal Regions	6
18	Shearlets	3
19	Gabor	7

Table II is the notation used hereafter.

TABLE II.   LIST OF NOTATION.

$\mathcal{T}$	Ordered set of image transforms
T_i	Single image transform i
ϕ_i	Parameters for image transform i
N	Population size
τ	GA termination criteria
$\mathbf{x}_j$	Individual j
n_j	j^{th} individual's number of genes
p_c	Crossover probability
p_m	Mutation probability
Θ	Diversity promoting constraints
f	Feature descriptor
F_j	Fitness for individual j

As in the ECO framework, iECO allows individuals chromosome's be of variable length. Herein, a chromosome is the segment of genes (i.e., series of image transforms) that represents a potential solution to the optimization task at hand. It is important to note that the ordering of genes does matter. Dilating an image followed by a distance transform will produce a different output than if conducted in the reverse order. In [1], the maximum allowed chromosome length was limited to 8 and this convention is also used for this work. Furthermore, ECO allows sub-regions of the image to be learned (if desired, as stated in [1], this is not required). For this work, experiments are performed on a data set composed

of image chips that have been identified by a pre-screener as regions of interest (we refer the reader to [2], [3] for details). The image chips are sub-regions of the original image, this is one factor that lead us to designing iECO to be performed on the "entire" image. Thus, the resultant transformed image produced by an individual will be the same size as the original image chip. Adaboost and perceptron are implemented in [1] (and herein) to combine chromosomes and formulate a fitness score for each individual. As in [1], the fitness score is computed as

$$F = \frac{tp * 500}{tp + fn} + \frac{tn * 500}{tn + fp}, \quad (1)$$

where tp, tn, fp, and fn is the number of true positives, true negatives, false positives, and false negatives, respectively. Therefore, the fitness score will be an integer value in the range $[0, 1000]$, with higher values being better.

A. Unique Feature Descriptor Approach

The iECO framework first advances the ECO features framework by using the ECO pipeline as a preprocessing stage, followed by the use of well known feature descriptors. The benefits of doing this are many. For one, feature descriptors have been heavily studied by the computer vision community, and many are founded on robust statistical properties. Furthermore, feature descriptors can be more easily interpreted than the ECO features. This is desirable as it allows more knowledge to be gained about a given problem and potentially further advanced. Additionally, there are a number of methods put forth that preserve the spatial context in imagery, resulting in localization of the features extracted (e.g., deploying feature descriptors on each cell/patch of an image partitioned in a cell-structured configuration– see [4] for an example). We expect the ECO framework to be highly susceptible to data sets in which there is much variation. Meaning, the bulk of imagery in the data sets used in [1] are centered on the object of interest. The ECO features learn specific sub-regions in the imagery; what happens when a new instance comes in where the object of interest is in the upper-left corner of the image, i.e., not exactly where the ECO feature is focused. Our conjecture is that it would fail miserably under such scenarios.

Beyond replacing ECO features with feature descriptors, and perhaps the more novel aspect of this method, is the ideology that different feature descriptors require different transforms to enable their full potential in extracting pertinent information. Herein, we propose that one should learn the series of image transforms, $\mathcal{T}$, that gives a particular feature descriptor the best chance at capturing discriminative information for the given domain's problem. That is, different feature descriptors represent information differently. As a result, each requires a unique $\mathcal{T}$ such that the data is best presented to the descriptor for it to exploit its unique method of extracting discriminative features. It is important to emphasize that the learning of $\mathcal{T}$ and their parameters is done for *each* feature descriptor being used. Algorithm 1 summarizes this process.

The need for such an approach is further supported when looking at the transformed imagery that was learned for the different feature descriptors used herein. First, we briefly mention these feature descriptors. Three different descriptors were

Algorithm 1 Learn $\mathcal{T}$ for each feature descriptor.

```
 1: for each feature descriptor, f_i, do
 2:     Create and initialize population;
 3:     while stopping condition not true do
 4:         for each individual, x_j, j = (1, 2, ..., N) do
 5:             for each training image, I_k, do
 6:                 Process I_k with T_j;
 7:                 Compute f_i on transformed image;
 8:             end for
 9:             Train perceptron;
10:             for each image in holding set do
11:                 Process image with T_j;
12:                 Compute f_i on transformed image;
13:             end for
14:             Compute and assign fitness score for x_j;
15:         end for
16:     end while
17: end for
```

implemented: *histogram of oriented gradients* (HOG) [5], [6], [7], *edge histogram descriptor* (EHD) [8], [9], and a statistical-based descriptor (SD). The SD is a simple descriptor we put forth and is composed of the following: local mean, standard deviation, kurtosis, and L_2-Norm, and the difference between the local values and their corresponding global value. Here, "local" refers to each cell of the cell-structured configuration.

Figure 1 provides an example of each feature descriptor's resulting transformed image for an instance of the data set in which a target, i.e., object of interest, is present. In Fig. 1, it is quite obvious that each feature descriptor has a unique $\mathcal{T}$ that is preferred by the descriptors for extracting features on this domain. It is important to stress that the learned $\mathcal{T}$ is *problem domain dependent*. That is, when moving from one problem domain to the next, $\mathcal{T}$ should be re-learned so that it is optimized for that domain.
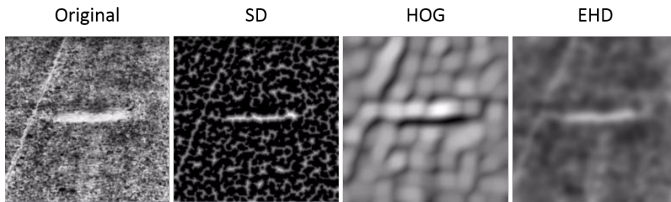

Fig. 1. Visualization of each feature descriptor's learned $\mathcal{T}$. From left to right: the original, $\mathcal{T}_{SD}$, $\mathcal{T}_{HOG}$, and $\mathcal{T}_{EHD}$ images are shown. Individually scaled for visual display.

At this point, we have discussed the portion of our iECO features framework that advances ECO features in two ways. One, feature descriptors are used instead of the ECO features, which (ECO features) are an unrolled image patch that has undergone a series of learned image transforms. And second, we propose a novel ideology to approaching feature descriptors for object recognition: condition the original data (through learned image transforms) for each feature descriptor. This allows each descriptor to extract better features, i.e., they are more discriminative. Now, we will discuss our approach to

ensuring that a diverse population is maintained throughout the lifetime of the GA through the implementation of constraints.

B. GA: Diversity Promoting Constraints

The importance of GAs having a diverse population has been well documented by the Evolutionary Computing community [10], [11], [12], [13]. First and foremost, a diverse population indicates that much of the search space is being explored. It is rather intuitive to realize that less diversity within a population will result in more regions of the search space being neglected. In standard GAs, mutation, which has a (typically low) probability of occurring, is the component that aims to add diversity to a population. However, the ECO search space is a rather unique and large optimization problem. Mutation alone is very unlikely to result in an adequate exploration by the population, i.e., pre-mature convergence to a sub-optimal solution is highly probable. As a result, a sophisticated approach needs to be incorporated into the GA for the search space to be more sufficiently explored.

Herein, we introduce diversity promoting constraints that attempt to consider the uniqueness and complexity of the ECO's search space. Two main questions were considered when designing constraints for this problem.

1) How much gene overlap allowed between individuals?: Further complicating this problem, the solution must be mindful of the path traversed, i.e., the ordering of the genes used that composes the individual. An illustration is provided in Fig. 2 to help with understanding. Additionally, there should be some amount of gene overlap allowed within the population so as to give regions of the space a fair chance to be adequately explored. This is illustrated in Fig. 3.
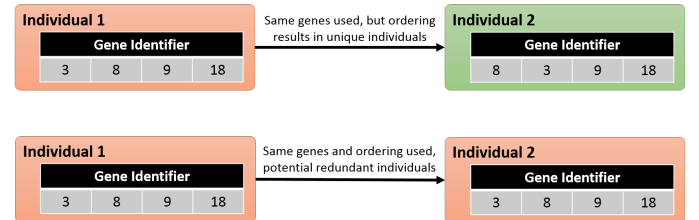

Fig. 2. Illustrating the reason that the diversity promoting constraints need to account for gene ordering.
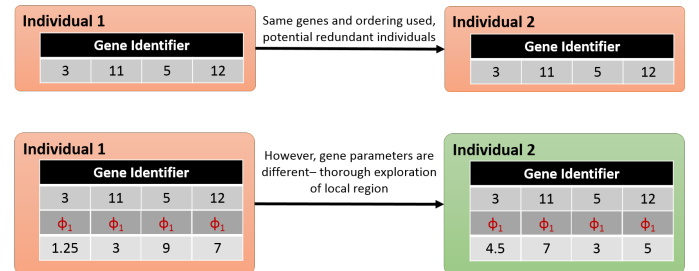

Fig. 3. Showing the need to allow some amount of gene overlap to occur.

2) How to address consecutive uses of the same gene?: It is our conjecture that the same gene occurring back-to-back is unwanted, an inefficient search of the same space (i.e., an unnecessary gene). Could the two genes be collapsed into one

by removal of the latter occurrence? Perhaps some middle-ground would be better, such as taking the mean between the repeated genes parameter values.

Considering the questions above, we propose the following solution. Designing a set of diversity promoting constraints, Θ that define what percentage of the population is allowed over-lapping genes at each layer of the individual's gene segment. For example, say that the maximum number of genes any one individual can have is 4. Then Θ would be a vector of length four, with the value at each index representing the population percentage that is allowed to overlap at the corresponding gene layer. It is important to understand that the constraint at the i^{th} layer of the gene segment is with respect to the sub-population in which all gene types (i.e., same image transform, not required to also have the same parameters) leading up to the i^{th} layer are the same. Figure 4 provides a graphical illustration to help with the understanding of this concept.

Next, we address the issue of having the same gene occur back-to-back in $\mathcal{T}$. As noted above, we believe that such a scenario is undesirable for most applications of this framework. For example, what sense does it make to take a rank transform twice, one right after the other? Additionally, repetitive genes in this manner increases the computational complexity of the system as a consequence of the unnecessary image transforms. In this work, we combat this by collapsing down consecutive uses of the same gene type. That is, if any particular gene occurs two or more times, consecutively, only the first occurrence is retained and the remaining consecutive genes of that type are removed. This results in the iECO framework also having the benefit of improving the efficiency of what is already a computationally expensive (ECO) system.

III. EXPERIMENTS

To assess the impact of the proposed iECO features frame-work, a comparative analysis between it and the original ECO features framework is given. This is done in the following two ways: 1) population diversity (Section III-A), and 2) system performance expressed by *Receiver Operating Characteristic* (ROC) curves (Section III-C). Additionally, we show that learning a unique $\mathcal{T}$ for each descriptor is both beneficial and necessary (Section III-B). Experiments are performed using the following set of parameters (which were empirically defined, but will be the subject of future work):

τ: 50 iterations
N: 50
max n: 8
p_c: 95%
p_m: 2%
Θ: $[0.15, 0.9, 0.8, 0.9, 0.8, 0.9, 0.7, 0.0]$

A. Experimental Analysis: Population Diversity

Table III is provided to help the reader realize how little diversity is achieved under the original ECO features framework, and is interpreted as follows. In the same manner as that in Fig. 4, as one traverses the layer's of the gene segment, the population % at the i^{th} layer is with respect to the sub-population from the preceding layer. All three descriptors' populations had more than 90% of its final population exploring relatively similar regions of the search space (i.e., same start gene). Of

TABLE III.   ECO FEATURES FRAMEWORK HAS MAJOR SHORTCOMING IN THE LACK OF POPULATION DIVERSITY.

| Population | Traversing Gene Segment: % Overlap Across Population | | | | | | | |
| | Gene Layer | | | | | | | |
	1	2	3	4	5	6	7	8
HOG	92%	78%	56%	60%	58%	71%	100%	100%
SD	92%	74%	56%	74%	50%	86%	83%	100%
EHD	96%	73%	49%	76%	35%	100%	50%	0%

the population with the same start gene, the same phenomenon exists throughout the entire length of the gene segment (as shown in Table III). There is indeed a need for a method to not only promote, but also ensure population diversity. Such a method is given in the iECO features framework in the form of Θ. It would be a moot point to give a similar table for the iECO features framework as it is guaranteed to abide by the constraints enforced by Θ.

Next, we look at the gene segments of the individuals representing the top 20% of each population at the completion of the learning algorithm for both the ECO and iECO features frameworks. This, along with each individual's fitness score is given in Tables IV and V for the ECO and iECO features framework, respectively. We only show this for the EHD feature descriptor; however, the same phenomenon holds true for the other two descriptors used herein. It is imperative to understand that the top 10% of the population are the elitist, therefore, there was no alteration of their genes. This is why most, if not all, of the first five individuals for each descriptor has gene segments that lack diversity– showing a sign of convergence. However, the next five individuals, from the iECO framework (Table V), have been subjected to the diversity promoting constraints and this is reflected in their gene segments. It becomes very obvious that iECO features framework results in a much more diverse population than the ECO features framework and thus, will have performed a more thorough exploration of the search space. This is a highly desirable property for GAs to possess as they are very susceptible to pre-mature convergence.

TABLE IV.   ECO FEATURES FRAMEWORK - f_{EHD}: TOP 10 INDIVIDUALS.

| Individual ID | Gene Layer | | | | | | | | Fitness |
	1	2	3	4	5	6	7	8	
1	18	8	8	7	13	7	-	-	840
2	18	8	8	7	13	7	-	-	840
3	18	8	8	7	13	7	-	-	840
4	18	8	8	7	13	7	8	7	840
5	18	8	8	7	13	7	7	-	840
6	18	8	7	8	7	8	8	7	840
7	18	8	8	7	13	8	7	8	840
8	18	8	7	7	8	7	7	-	840
9	18	8	7	8	8	7	13	18	840
10	18	7	-	-	-	-	-	-	840

Furthermore, Table IV provides a great example of how the ECO features framework allows for nonsensical ordering of genes. To elaborate, look at the first individual's second, third, and fourth gene. The second and third genes are consecutive *Log* image transforms, which has no parameter. This is followed by a rank transform– why the need for the third gene? The repeated use of the third gene has no impact on the
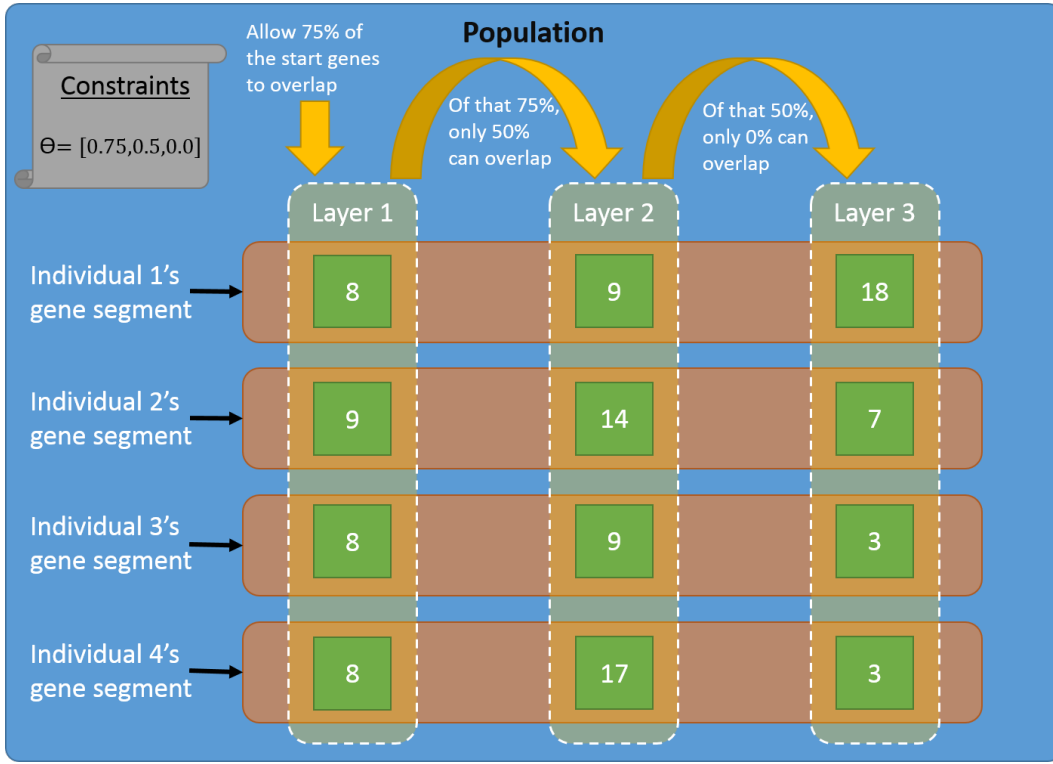
Fig. 4. Illustration of the implementation of diversity promoting constraints. Based on Θ, its first value, 0.75, ensures that no more than 75% of the population has the same start gene (represented by Layer 1). Its next value of 0.5 translates to the following. At most, 50% of the population with the *same* start gene can have the same second gene (represented by Layer 2). Finally, 0% ensures that no two individuals have the exact same chromosome.

TABLE V. iECO FEATURES FRAMEWORK - f_{EHD}: TOP 10 INDIVIDUALS.

| Individual ID | Gene Layer | | | | | | | | Fitness |
	1	2	3	4	5	6	7	8	
1	3	11	5	-	-	-	-	-	886
2	3	11	5	-	-	-	-	-	886
3	3	11	5	-	-	-	-	-	886
4	3	11	5	-	-	-	-	-	886
5	3	11	5	-	-	-	-	-	886
6	8	5	11	5	4	5	11	-	886
7	4	5	3	11	5	-	-	-	863
8	1	11	5	-	-	-	-	-	863
9	13	3	11	5	8	11	5	-	818
10	12	11	5	-	-	-	-	-	818

fourth gene's result. Such phenomena occurred consistently in all of the experiments conducted using the ECO features framework (including those not reported herein). Also, for this problem at least, the necessity for using feature descriptors over unrolling an image patch is realized as eight of the top ten iECO individuals result in higher fitness scores than does the best ECO individual. Lastly, comparing Tables IV and V, it is obvious that the iECO framework tends to produce individuals of shorter length than does the ECO approach (i.e., better computational efficiency).

B. Experimental Analysis: Unique $\mathcal{T}$ Impact

In Section II-A, it was proposed that a unique $\mathcal{T}$ should be learned for each feature descriptor being used. Additionally, it was shown in Fig. 1 that each descriptor does, visually,

appear to be learning something entirely different. We now give a more analytical approach to supporting the need for such an approach. To do this, we computed the fitness score for the top 10 individuals from each population, but having them extract the other feature descriptors used herein. Their resulting fitness score and corresponding % change (Δ) over using its own feature descriptor is given in Tables VI-VIII.

TABLE VI. EXTRACT HOG USING TOP INDIVIDUALS FROM SD AND EHD POPULATIONS.

| Individual ID | SD | | EHD | |
	F	% Δ	F	% Δ
1	818	-3%	818	-3%
2	818	-3%	795	-5%
3	818	-3%	795	-5%
4	818	-3%	795	-5%
5	818	-3%	795	-5%
6	795	-5%	795	-5%
7	795	-5%	795	-5%
8	795	-5%	773	-8%
9	795	-5%	773	-8%
10	795	-5%	773	-8%

In all instances, the fitness score dropped. From this, we can infer that each descriptor's learned $\mathcal{T}$ is indeed unique and beneficial to that descriptor.

C. Experimental Analysis: System Performance

We now present results for implementing the iECO framework on a real-world data set with application for automatic detection of buried explosive hazards in forward looking-long

TABLE VII. EXTRACT SD USING TOP INDIVIDUALS FROM HOG AND EHD POPULATIONS.

Individual ID	HOG		EHD	
	F	% Δ	F	% Δ
1	705	-18%	773	-11%
2	705	-18%	773	-11%
3	705	-18%	773	-11%
4	705	-18%	750	-13%
5	705	-18%	727	-16%
6	705	-18%	682	-21%
7	705	-18%	659	-24%
8	682	-21%	659	-24%
9	682	-21%	659	-24%
10	659	-24%	659	-24%

TABLE VIII. EXTRACT EHD USING TOP INDIVIDUALS FROM HOG AND SD POPULATIONS.

| Individual ID	HOG		SD	
	F	% Δ	F	% Δ
1	818	-8%	795	-10%
2	750	-15%	773	-13%
3	727	-18%	773	-13%
4	727	-18%	773	-13%
5	727	-18%	750	-15%
6	705	-21%	750	-15%
7	705	-21%	727	-18%
8	705	-21%	727	-18%
9	705	-21%	705	-21%
10	682	-23%	705	-21%

wave infrared (FL-LWIR) imagery. This data set consists of 16 runs across two lanes at an arid United States Army test site. Of the 16 runs, 7 were from lane 1 and 9 from lane 2. Targets are buried at different depths, ranging from 1 to 6 inches, and also vary in terms of metal content, i.e., heavy to low to no metal content. Additionally, this data set was collected during the morning and afternoon to include the thermal variations that occur at different times of the day. Lane-based cross validation (CV) was used for testing, e.g., train an SVM using runs from lane 1 and test using runs from lane 2. Because co-plotting multiple ROCs on the same figure can be very difficult to view, we vertically average each individual runs ROC belonging to the lane being tested.

To assess the impact of the iECO framework, we compare its results with the ECO framework in two ways. First, a direct comparison between iECO and ECO are given. As will be shown, iECO vastly outperforms ECO features in this setting. Second, we show that the iECO framework also outperforms the ECO framework even if using the feature descriptors instead of ECO features (i.e., unrolling the transformed image). We believe that this is likely a result/indicative of the ECO framework's poor searching of the solution space. Learning for both methods was performed on a very small subset of the training data (roughly 1%), half of which were targets (class 1) and half non-target (class 2). Lastly, results are given for experiments using only the best individual as well as using the top 5 individuals. This was done to explore the idea that the different individuals could each be learning something independent from each other that, when combined, improves system performance. To ensure clarity, the top y individuals from *each* descriptor is used, not the globally best y individuals and their corresponding descriptors.

1) iECO Outperforms: First, an assessment of the need for using feature descriptors instead of ECO features is given. Figure 5 shows the vertically averaged ROC curves along with the 95% confidence intervals resulting from the two approaches.
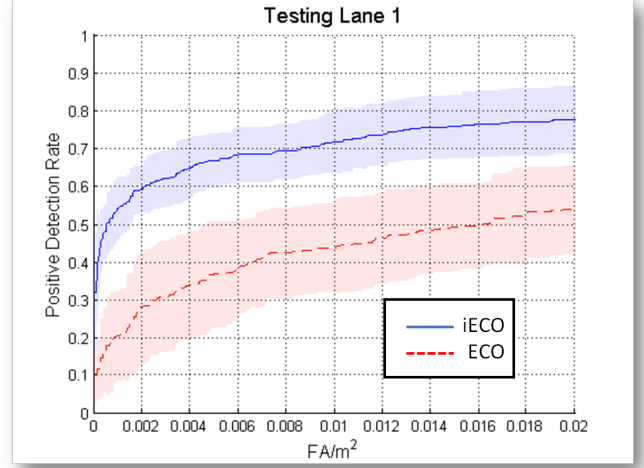
(a) Testing lane 1
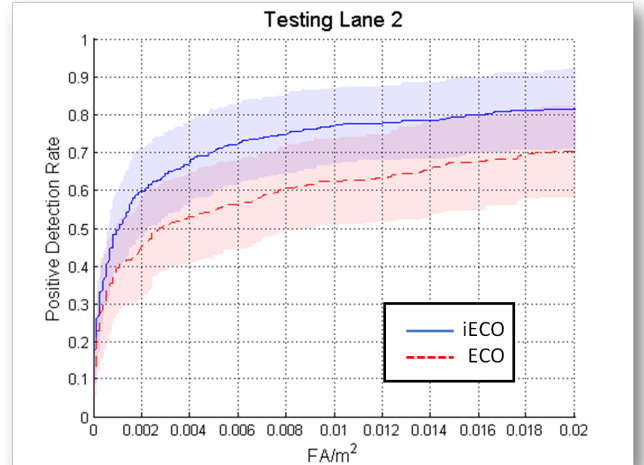
(b) Testing lane 2

Fig. 5. Vertically averaged ROC curves with 95% confidence intervals. iECO is shown as the solid blue line, and ECO as the dashed red line. For each, 95% confidence interval is overlaid in corresponding color.

Here we see iECO drastically outperforms ECO features. This supports our hypothesis that using cell-structured feature descriptors provides a much more robust method for capturing pertinent image data than does the ECO features. Additionally, this data set's vast collection of varying target signatures (i.e., size, shape, texture, etc.) potentially exploits the ECO features vulnerability to intra-class variation. That is, the ECO features' need for objects of interest to be presented in a relatively static/repetitive setting (e.g., images of faces that are centered and the focus of the imagery).

From this, the following question may arise. A major part of this paper's contribution is in the diversity promoting constraints, how do we know that they are contributing to per-

formance improvement?– Is improvement simply the result of using feature descriptors on the $\mathcal{T}$ transformed imagery instead of ECO features? To investigate such a scenario, experiments were performed using the ECO framework; however feature descriptors were applied to ECO's $\mathcal{T}$ transformed imagery.

2) iECO– Is it Just the Descriptors?: We now provide preliminary results on the investigation into the impact iECO's diversity promoting constraints has on a real-world application. It was already shown that Θ does have a positive impact in regards to ensuring population diversity is sustained, and thus, a more thorough search of the space is achieved. However, how does this translate in terms of system performance. Figures 6 and 7 show the vertically averaged ROC curves and 95% confidence intervals for experiments using the best individual and the top 5 individuals from each descriptors population, respectively.
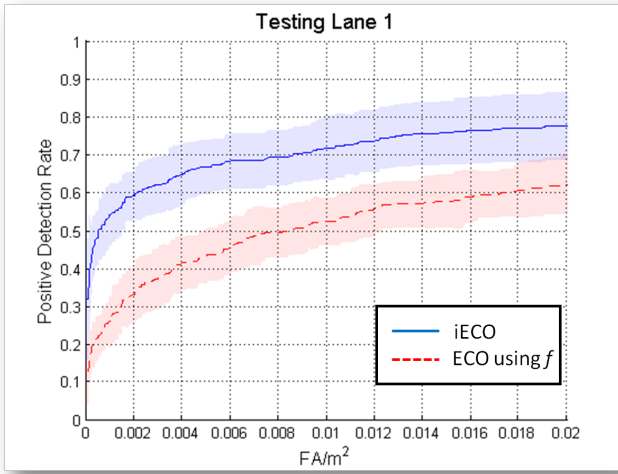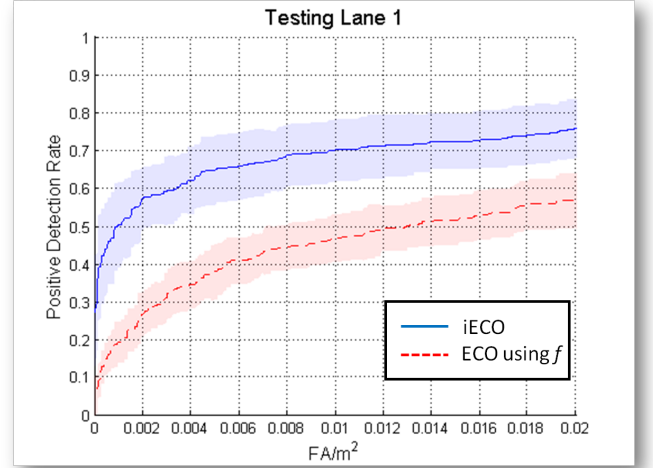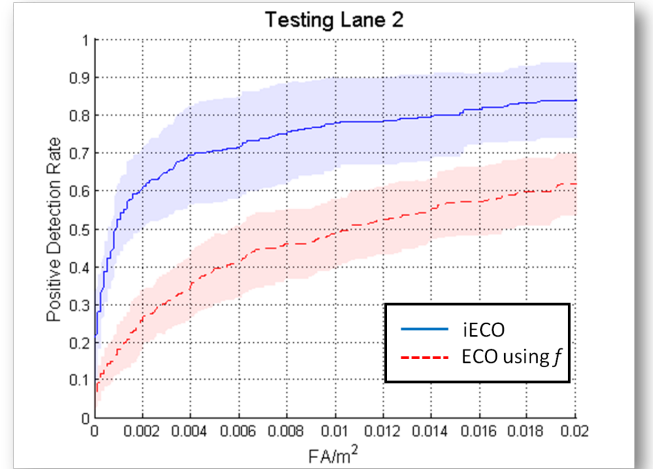

(a) Testing lane 1
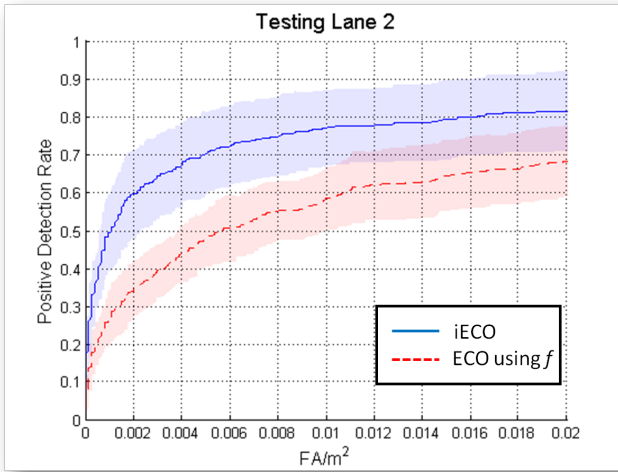
(b) Testing lane 2

Fig. 6. Vertically averaged ROC curves with 95% confidence intervals produced using each f's best individual. iECO is shown as the solid blue line, and ECO as the dashed red line. For each, 95% confidence interval is overlaid in corresponding color.

Interestingly, iECO largely outperforms the ECO using

(a) Testing lane 1

(b) Testing lane 2

Fig. 7. Vertically averaged ROC curves with 95% confidence intervals produced using each f's top 5 individuals. iECO is shown as the solid blue line, and ECO as the dashed red line. For each, 95% confidence interval is overlaid in corresponding color.

feature descriptors approach. Noting that the only difference between the two methods in iECO and ECO using feature descriptors (really, a hybrid of ECO and iECO) is that iECO includes the diversity promoting constraints. Seeing that ECO using f (best individual from each only– Fig 6) does in fact perform better than the ECO features reported in 5, support is gained for the introduced Θ being an key contributor to the overall improvement in system performance.

Lastly, a potentially unexpected result comes from comparing iECO's performance in Figs. 6 and 7. That is, performance is actually better if only the best individual from each f is used instead of a collection of top individuals. We do remind the reader that a very simply and straight-forward approach to feature fusion was used herein (feature vector concatenation)– future work is planned for investigating more sophisticated feature level fusion, which is likely required to see the full potential of the iECO framework.

IV. Conclusion

Herein, we proposed the iECO features framework. This advances the ECO features by using feature descriptors rather than a simple unrolling of a transformed image patch and also through the introduction of diversity promoting constraints. It was shown that iECO has a number of benefits and improvements to the ECO framework. Experimental results suggest that our constraints lead to a better searching of the solution space. Additionally, iECO tends to find individuals whose chromosomes are smaller in length, and therefore less computationally complex then the ECO algorithm. Furthermore, feature descriptors were shown to provide great performance improvement over ECO features. In future work, we plan to look into developing novel fusion techniques for improving the ability to identify discriminatory information for the different individuals and across the different feature descriptors. Additionally, we are developing methods that will track which volumes of the search space have been explored and ensure that the population is thoroughly diverse.

References

[1] K. Lillywhite, D.-J. Lee, B. Tippetts, and J. Archibald, "A feature construction method for general object recognition," *Pattern Recognition*, vol. 46, no. 12, pp. 3300–3314, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0031320313002549

[2] K. Stone, J. Keller, D. Anderson, and D. Barclay, "An automatic detection system for buried explosive hazards in fl-lwir and fl-gpr data," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2012, pp. 83 571E–83 571E.

[3] K. Stone, J. M. Keller, M. Popescu, and C. J. Spain, "Buried explosive hazard detection using forward-looking long-wave infrared imagery," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2011, pp. 801 725–801 725.

[4] S. R. Price, D. T. Anderson, K. Stone, and J. M. Keller, "Investigation of context, soft spatial and spatial-frequency domain features for buried explosive hazard detection in fl-lwir," in *Proc. SPIE 9072, Detection and Sensing of Mines, Explosive Opjects, and Obscured Targets*, 2014, pp. 907 217–907 217–14. [Online]. Available: http://dx.doi.org/10.1117/12.2049937

[5] S. Edelman, N. Intrator, and T. Poggio, "Complex cells and object recognition," 1997.

[6] D. G. Lowe, "Object recognition from local scale-invariant features," in *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, vol. 2. Ieee, 1999, pp. 1150–1157.

[7] ——, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.

[8] H. Frigui and P. Gader, "Detection and discrimination of land mines in ground-penetrating radar based on edge histogram descriptors and a possibilistic-nearest neighbor classifier," *Fuzzy Systems, IEEE Transactions on*, vol. 17, no. 1, pp. 185–199, 2009.

[9] K. Stone, J. Keller, D. Anderson, and D. Barclay, "An automatic detection system for buried explosive hazards in fl-lwir and fl-gpr data," in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2012, pp. 83 571E–83 571E.

[10] E. Burke, S. Gustafson, and G. Kendall, "Diversity in genetic programming: an analysis of measures and correlation with fitness," *Evolutionary Computation, IEEE Transactions on*, vol. 8, no. 1, pp. 47–62, Feb 2004.

[11] E. Zitzler, K. Deb, and L. Thiele, "Comparison of multiobjective evolutionary algorithms: Empirical results," *Evolutionary computation*, vol. 8, no. 2, pp. 173–195, 2000.

[12] J. Arabas, Z. Michalewicz, and J. Mulawka, "Gavaps-a genetic algorithm with varying population size," in *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*. IEEE, 1994, pp. 73–78.

[13] R. K. Ursem, "Multinational evolutionary algorithms," in *Proceedings of the Congress on Evolutionary Computation*. IEEE Press, 1999, pp. 1633–1640.